The Dune FoamGrid implementation for surface and network grids

Oliver Sander[1], Timo Koch[2], Natalie Schröder[2], and Bernd Flemisch[2]

[1]TU Dresden, Institute for Numerical Mathematics, oliver.sander@tu-dresden.de

[2]University of Stuttgart, Institute for Modelling Hydraulic and Environmental Systems,

{timo.koch, natalie.schroeder, bernd.flemisch}@iws.uni-stuttgart.de

arXiv:1511.03415v1 [cs.MS] 11 Nov 2015

November 12, 2015

We present FoamGrid, a new implementation of the Dune grid interface. FoamGrid implements one- and two-dimensional grids in a physical space of arbitrary dimension, which allows for grids for curved domains. Even more, the grids are not expected to have a manifold structure, i.e., more than two elements can share a common facet. This makes FoamGrid the grid data structure of choice for simulating structures such as foams, discrete fracture networks, or network flow problems. FoamGrid implements adaptive non-conforming refinement with element parametrizations. As an additional feature it allows removal and addition of elements in an existing grid, which makes FoamGrid suitable for network growth problems. We show how to use FoamGrid, with particular attention to the extensions of the grid interface needed to handle non-manifold topology and grid growth. Three numerical examples demonstrate the possibilities offered by FoamGrid.

1 Introduction

Various simulation problems are posed on domains that are not open subsets of a Euclidean space. Frequently, such domains are surfaces or curves embedded in a higher-dimensional Euclidean space. Equations on such domains, sometimes called geometric partial differential equations, comprise diffusion and transport on the surface [Dziuk and Elliott, 2007b], flow problems [Nitschke et al., 2012, Reuther and Voigt, 2015], and phase-field equations [T. Witkowski, 2012]. Sometimes, movement of the surface itself is modeled [Dziuk and Elliott, 2007a], and this movement may couple with processes on the surface [Gross and Reusken, 2011].

As an additional difficulty, some boundary value problems are posed on domains Ω that do not even have the structure of a topological manifold. That is, not every point of Ω has a neighborhood that is homeomorphic to an open subset of $\mathbb{R}^d$. Figure 1 illustrates this: While the surface patch on the left is locally homeomorphic to Euclidean space, the one in the middle is a T-junction, and the one on the right is a touching point. Two-dimensional domains with such features appear in applications like the simulation of closed-cell foams [Nammi et al., 2010], or networks of fractures in rock mechanics [McClure and Horne, 2013, McClure et al., 2015]. Of considerable importance are also one-dimensional networks embedded into a two- or three-dimensional Euclidean space. They appear in models of traffic networks [Garavello and Piccoli, 2006], supply chains [D'Apice et al., 2010], but also in simulations of biological systems like root networks [Dunbabin et al.,

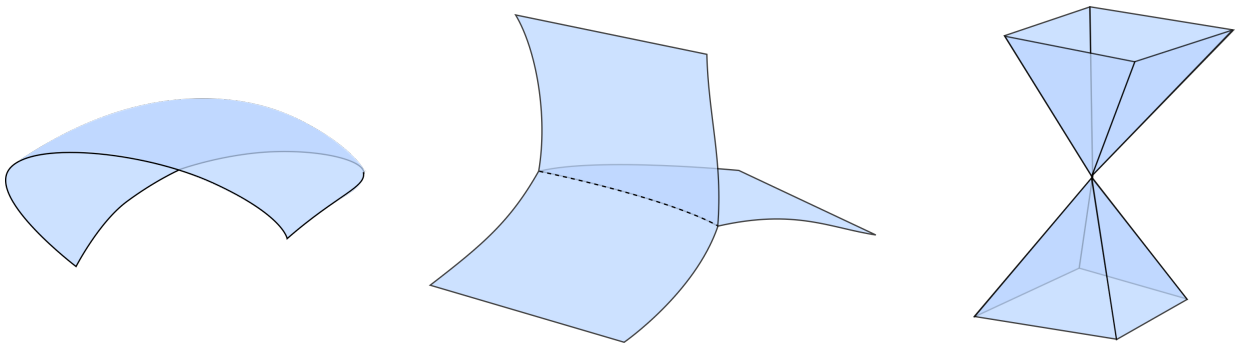

Figure 1: Computational domains might not be topological manifolds. Left: manifold; center: T-junction; right: touching point

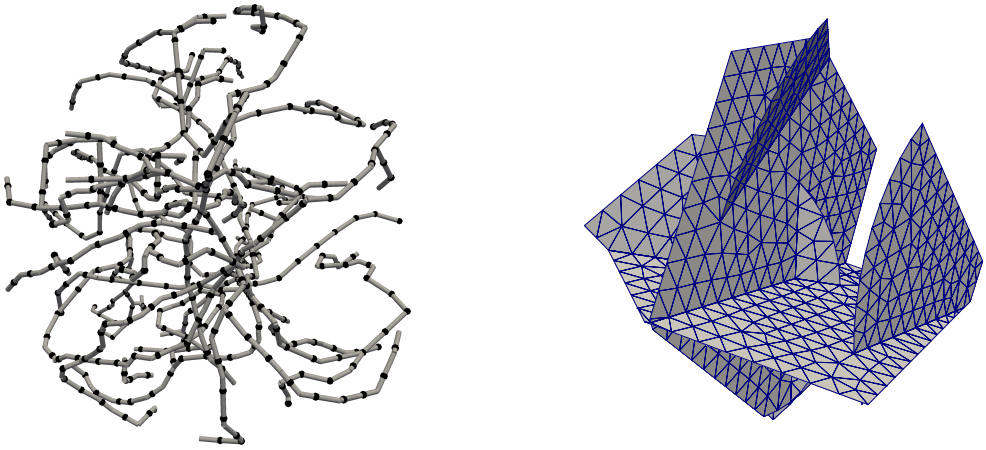

Figure 2: One- and two-dimensional network grids

2013], neural networks [Lang et al., 2011], or blood vessel networks [Cattaneo and Zunino, 2014a]. An overview over flow problems on networks is given in [Bressan et al., 2014]. Figure 2 shows two example domains for network problems.

Various discretization methods have been proposed for surface and network equations [Dziuk and Elliott, 2013, Olshanskii et al., 2009]. Explicit discretizations, which are the focus of this work, use a grid of the same dimension as the domain. For manifold surface grids it is reasonably simple to generalize grid data structures to such a setting. The main hurdles are admitting that the number of coordinates of a vertex can be different from the effective grid dimension, and making sure that grids are not required to have a boundary. Several standard simulation codes support such surface grids. We mention Alberta [Schmidt and Siebert, 2005] (standalone and as part of Dune), AMDiS [Vey and Voigt, 2007], and FEniCS [Rognes et al., 2013]. Moreover, the GeometryGrid Dune meta grid allows to embed any Dune grid into a Euclidean space of higher dimension.

Grid data structures for non-manifold grids are more challenging. To handle T-junctions, for example, the data structure must cope with the fact that element facets (i.e., edges in a 2d grid) may have more than two neighbors (Figure 1). While this is not very difficult to implement, it requires the introduction of additional data fields and logic, which is not used when the grid happens to be a manifold. Since the latter case is predominant, such additional features mean space and run-time overhead for most users. Therefore, standard grid data structures do not allow for non-manifold topologies. Numerical examples of processes on non-manifold topologies in the

literature are typically done using ad hoc implementations of the necessary grid data structures.

Unfortunately, using ad hoc implementations wastes a lot of human resources. While the domain may be non-standard, many of the equations on network grids differ little from their Euclidean counterparts. However, existing code like finite element assemblers for diffusion and transport processes cannot be reused on ad hoc grid data structures. They require detailed knowledge of the grid, and it is a challenge to port an existing assembler to a new grid data structure. Also, since the data structure is ad hoc, it is difficult to reuse it in other contexts. In particular, it is difficult to share it with other groups working in the same field.

The Dune software system[1] [Bastian et al., 2008b,a] has found an elegant solution for both problems. Dune is a set of open-source C++ libraries dedicated to various aspects of finite element and finite volume methods. Its grid component, implemented in the `dune-grid` module, specifies an abstract interface for computational grids. The specification mandates what a data structure should be able to do to qualify as a Dune grid, and how this functionality should be accessible from C++ code. Examples of such functionality include being able to iterate over the elements and vertices, and getting the maps from the reference element to the grid elements. Those parts of a numerical simulation code that use the grid, such as the matrix assemblers or error estimators, are written to use only the abstract grid interface. Grid data structures implementing this interface are then completely decoupled from the algorithms that use them.

This decoupling has various interesting consequences. Using the Dune grid interface, it is easy to swap a given grid implementation for another one. Indeed, in typical Dune applications the C++ grid type is set once in the code, and handed around as a template parameter. Changing the initial `typedef` and recompiling the code is usually sufficient to switch to an alternative grid data structure. This possibility to easily switch between grid implementations allows to provide tailor-made grid data structures for special simulation needs. For example, `dune-grid` itself provides YaspGrid, the implementation of a structured grid with very little run-time and space overhead. In contrast, UGGrid implements a very flexible unstructured grid with non-conforming and red–green refinement.

In this paper we present FoamGrid, a new implementation of the Dune grid interface that is dedicated to surface and network grids. Its main features are:

- FoamGrid implements one- and two-dimensional simplex grids embedded in a physical Euclidean space of arbitrary dimension w. Hence, it targets geometric and surface PDEs.
- The grids do not need to have the structure of a topological manifold. Network configurations like the ones in Figures 1 and 2 are supported.
- A FoamGrid can be adaptively refined. In the standard setup, refining a triangle results in four coplanar triangles. Additionally, FoamGrid elements can be parameterized, i.e., they can be given a map that describes a nonlinear embedding of the element into $\mathbb{R}^w$. As the element gets refined more and more, its shape approaches the one described by the parameterization (Figure 4).
- Finally, the domain of a FoamGrid can grow and shrink at run-time. Elements can be added and removed even when there is no coarser "father" element, without invalidating the grid data structure. Data can be transferred during this process. This allows to elegantly simulate network growth and remodeling processes.

FoamGrid is available in a Dune module `dune-foamgrid`,[2] and is installed just like any other Dune module. `dune-foamgrid` is free software, available under the either the LGPLv3+, or the GPLv2 with a linking exception clause.

[1] `www.dune-project.org`

[2] `http://users.dune-project.org/projects/dune-foamgrid`

Using Dune and FoamGrid for simulations of network and surface PDE problems has a number of important advantages. First of all, you do not have to implement the grid data structure yourself. While not overly difficult, quite a bit of thought has gone into FoamGrid, which would need work to be replicated. Then, since FoamGrid implements the Dune grid interface, large amounts of existing application and infrastructure code can be used directly with FoamGrid. This includes things like finite element spaces and assemblers, error estimators, and grid file readers and writers. This advantage is demonstrated in particular by the numerical example in Section 4.1, which required no additional coding at all to extend an existing planar code to a network setting.

As a further advantage, once a user starts employing FoamGrid and the Dune grid interface, he immediately has the power of all other Dune grid implementations at his disposal. All of these are easily usable together with FoamGrid. Hence, e.g., in simulations that couple network grids with background grids, several implementations of such background grids are readily available.[3] Since the Dune grid interface is a well-established standard, it is easy to learn how to use FoamGrid. If a user already knows Dune, there is little additional knowledge needed. Since FoamGrid is open source, sharing code based upon it is particularly easy.

While FoamGrid has many interesting features, there is also a number of things it does not currently support. For example, elements can only be simplices, and they must be one- or two-dimensional. (The authors could not think of a use case for a higher-dimensional network grid. Write us if you know one.) Adaptive grid refinement is currently non-conforming, which leads to hanging nodes, and, possibly, to holes in the surface (Figure 5). Finally, the current FoamGrid implementation is purely sequential, and FoamGrid objects cannot be distributed across several processes. However, the development of FoamGrid is ongoing, and these features may appear in later releases.

The present article is structured as follows. Chapter 2 briefly explains how to use FoamGrid. Everything mentioned there is codified in the Dune grid interface specification, and hence you can also read this chapter as an introduction to the use of the Dune grid interface. Chapter 3 then explains the special features that FoamGrid offers. In particular, these are support for T-junctions, adaptive refinement with element parameterizations, and the ability to "grow". Finally, in Chapter 4 we give three numerical examples showcasing the different features of FoamGrid. The first shows unsaturated flow through a two-dimensional fracture network using finite elements. The second example shows h-adaptive, locally mass-conservative transport of a therapeutic agent in a microvascular network using finite volumes. The third one models the growth of plant root networks.

2 FoamGrid and the Dune grid interface

In this chapter, we start by describing the programmer interface of FoamGrid. FoamGrid implements the Dune grid interface, hence in many central aspects, it can be used just like any other Dune grid. This chapter focuses on these aspects. You can therefore also read it as a brief review of the Dune grid interface. For more details, you may want to consult the Dune online documentation and [Bastian et al., 2008b,a].

The central class of the FoamGrid grid implementation is

```
1 template <int dim, int dimworld>
2 class FoamGrid;
```

available from the header `dune/foamgrid/foamgrid.hh`. This class implements a hierarchical grid as defined in [Bastian et al., 2008b, Def. 13], i.e., a coarse (or *macro*) grid, and element refinement trees rooted in each of its elements. The first template parameter `dim` is the grid dimension d,

[3] And the Dune extension module `dune-grid-glue` offers a convenient way to do the coupling (`www.dune-project.org/modules/dune-grid-glue`).

which must be either 1 or 2. The second template parameter `dimworld` is the dimension w of the Euclidean embedding space. It must be equal to or greater than the grid dimension, but can otherwise be arbitrary. For the rest of this article we use `dim` and `dimworld` in code examples, and d and w in text to denote the dimensions of the grid and the physical space, respectively.

To construct FoamGrid objects, the `FoamGrid` class implements the entire Dune grid interface for the setup of unstructured grids. In particular, the class `GridFactory` is implemented for `FoamGrid`. Thus, all file-reading methods based on this interface are available. For example, files in the gmsh format [Geuzaine and Remacle, 2015] can be read by using the line

```
std::shared_ptr<FoamGrid<2,3>  grid(
    GmshReader<FoamGrid<2,3>>::read("filename.msh") );
```

This will read the file named "`filename.msh`", and set up a new `FoamGrid<2,3>` object with it. The new grid object is returned in a shared pointer called `grid`. Note that vertex coordinates in gmsh files always have three components, so reading a gmsh file into a `FoamGrid<1,2>` object will discard the third entry. As a special feature, gmsh files can contain elements with polynomial geometries of order up to five. While FoamGrid element geometries are always affine, FoamGrid can use the higher-order geometries during mesh refinement (Section 3.2).

Writing `FoamGrid` objects to disk is equally straightforward. All Dune file writing codes rely on the grid interface only, and can therefore be used with FoamGrid. For example, writing the object pointed to by the `grid` shared pointer into a VTK file called `my_filename.vtu` can be achieved by including the header `dune/grid/io/file/vtk.hh` from the `dune-grid` module, and writing

```
typedef FoamGrid<2,3> GridType;
VTKWriter<GridType::LeafGridView> vtkWriter(grid->leafGridView());
vtkWriter.write("my_filename");
```

The resulting file can be visualized, e.g., with the ParaView software.

2.1 Elements and geometries

In Dune, finite element assembly typically takes place on the leaf elements of the refinement trees. These elements form the *leaf grid view*, which encapsulates the notion of a textbook non-hierarchical finite element grid. From a FoamGrid, the leaf grid view can be obtained in the usual way, i.e.,

```
auto foamGridLeafView = grid->leafGridView();
```

which already has been used in the VTK example above.

Access to grid elements and vertices is provided by the grid view. Elements can be iterated over using `begin`/`end`-iterators

```
for (auto it = foamGridLeafView.begin<0>();
     it != foamGridLeafView.end<0>();
     ++it)
{
  // do something with the element in '*it'
}
```

where the number `0` specifies that the loop is to be over the grid elements (it is the codimension of the grid elements with respect to the grid). Using the C++11 range-based-`for` syntax, the same loop can be written more concisely

```
for (const auto& element : elements(foamGridLeafView))
{
  // do something with the element in 'element'
}
```

Similarly, a loop over all vertices of the grid is written as

```
for (auto it = foamGridLeafView.begin<dim>();
     it != foamGridLeafView.end<dim>();
     ++it)
{
  // do something with the vertex in '*it'
}
```

In this code, the number `dim` (i.e., the grid dimension), specifies that the loop is to be over the grid vertices, because vertices are zero-dimensional and hence have codimension `dim` in a `dim`-dimensional grid. Alternatively, one can write

```
for (const auto& vertex : vertices(foamGridLeafView))
{
  // do something with the vertex in 'vertex'
}
```

The objects `vertex` and `element` (or `*it` in the iterator loops) are instances of what in Dune terminology are called *entities*. Entities are implemented by the `Entity` interface class in `dune-grid`. They provide topological information about the grid elements and vertices, like links to the corners vertices of an element, and to the father and descendant elements in the refinement tree. In all these aspects, `FoamGrid` objects behave just like any other Dune grids. Furthermore, each element entity provides a `Geometry` object, which represents the affine map $F : T_{\text{ref}} \to T \subset \mathbb{R}^w$ from the reference element T_{ref} to the grid element T. This map provides the geometrical information needed to assemble finite element and finite volume systems, like evaluation of F and its inverse F^{-1}, the inverse transposed Jacobian matrix $\nabla F^{-\mathrm{T}}$, and the functional determinant $J := \sqrt{\det \nabla F^{\mathrm{T}} \nabla F}$. Note that since for surface grids the world dimension w is larger than the dimension of the reference element T_{ref}, the image of T_{ref} under F is a set of measure zero in $\mathbb{R}^w$. In finite-precision arithmetic the argument of the method implementing F^{-1} will typically not be in the domain of F^{-1}. The `Geometry` implementation of FoamGrid therefore extends F^{-1} to the entire space $\mathbb{R}^w$. Given a point $x \in \mathbb{R}^w$ not necessarily in T, `FoamGrid` computes a point ξ in the plane spanned by T_{ref} such that $|F(\xi) - x|$ is minimized. The affine function F is tacitly extended from T_{ref} to its entire affine hull here.

To compute element fluxes, elements in a Dune grid provide a set of so-called *intersections*, which relate elements to their neighbors. This mechanism is very flexible, and in particular handles general non-conforming situations very well. Nevertheless, using grids for network domains stretches the bounds of the current intersection concept, and some generalization is needed. We have therefore dedicated a separate chapter to FoamGrid intersections (Chapter 3.1).

2.2 Attaching data to grids

Data is attached to FoamGrid objects in the same way as to other Dune grids. Each grid view object can provide a corresponding `IndexSet` object

```
const auto& indexSet = foamGridLeafView.indexSet();
```

This index set provides an integer number for each entity (i.e., vertex, edge, or element) of the grid view. For each dimension and reference element type, these numbers are consecutive and start at zero. They can hence be used to address random-access containers holding the simulation data. This approach is convenient, flexible, and efficient.

Data stored in arrays is lost if the grid changes, either by refinement (Section 2.3) or by grid growth (Section 3.3). To preserve data across grid modifications, FoamGrid, just like any other Dune grid, additionally provides a set of persistent numbers. These are obtained by an `IdSet` object

```
const auto& idSet = grid->localIdSet();
```

(remember that in our initial example the variable `grid` was a shared pointer to a `FoamGrid`). Persistent numbers are neither consecutive nor restricted to start at zero, but they can be used to access search trees or hash maps. Before modifying the grid, all simulation data must be copied into such data structures, and copied back to arrays after the modification is completed. While such copying is costly, its run-time is usually negligible compared to the cost of the actual grid modification.

2.3 Adaptive refinement

FoamGrid supports red refinement (non-conforming refinement) of simplices, where each triangle is split into four congruent smaller triangles. If a two-dimensional FoamGrid is refined locally, then hanging nodes appear in the grid. Depending on the discretization used, this may or may not be a problem. True red–green refinement, which avoids the hanging nodes, may appear in later versions of FoamGrid.

Adaptive grid refinement in FoamGrid is controlled via the standard Dune grid interface. In a first step the method

```
bool mark (int refCount, const Codim<0>::Entity& element);
```

is used to mark an element `element` for refinement (`refCount` > 0) or coarsening (`refCount` < 0). The method returns `true` if the element was successfully marked. The mark of an element `element` can be obtained with the method

```
int getMark (const Codim<0>::Entity& element) const;
```

which returns `1` for elements marked for refinement, `-1` for elements marked for coarsening, and `0` for unmarked elements.

The grid is then modified in a second step with the methods

```
bool coarsen = grid.preAdapt(); // true if at least one element
                                // will be coarsened
bool refined = grid.adapt();    // true if at least one element
                                // was refined
grid.postAdapt();
```

Between `preAdapt()` and `adapt()` it is possible to check the following flag

```
bool mightVanish = element.mightVanish(); // true if the element might
                                          // vanish due to coarsening
                                          // by grid.adapt()
```

Similarly, between `adapt()` and `postAdapt()` one can check the flag

```
bool isNew = element.isNew(); // true if element was created by last
                              // call to adapt()
```

Both are useful to manage the transfer of data associated with the grid. As FoamGrid itself does not store any associated data, the data transfer from the old grid to the adapted grid is managed by the user. An example featuring adaptive refinement and coarsening is presented in Section 4.2.

3 Additional features

The previous chapter has described those aspects where FoamGrid behaves just like any other Dune grid. However, FoamGrid also has a few features that set it apart from most other Dune grids. The present chapter is dedicated to those features.

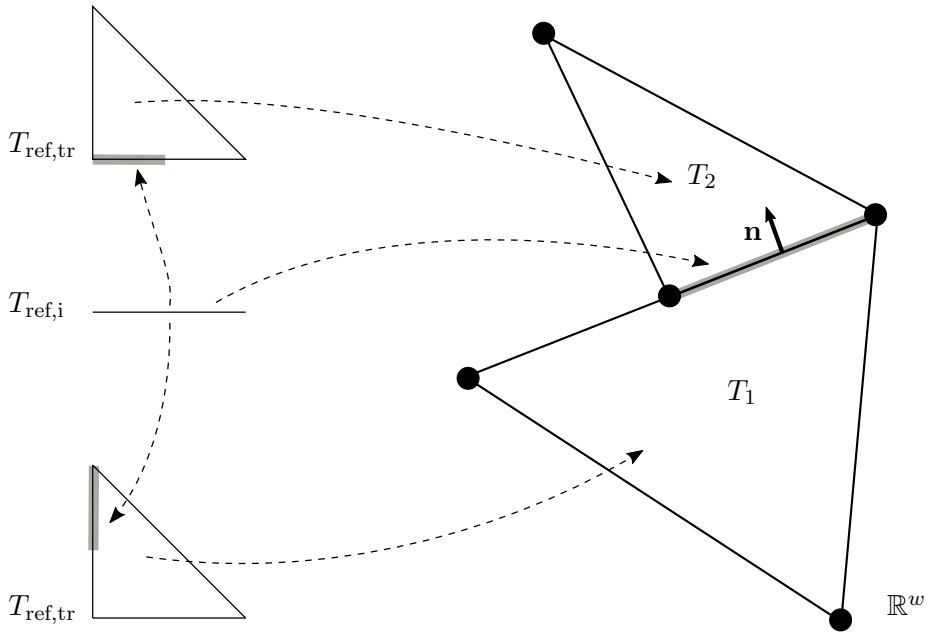

Figure 3: Intersection between two elements T_1 and T_2. From its reference element $T_{\text{ref,i}}$, there are two maps to $T_{\text{ref,tr}}$ (the reference element of T_1 and T_2), and one to $\mathbb{R}^w$, which describe the shape of the intersection in T_1, T_2, and $\mathbb{R}^w$, respectively.

3.1 Handling intersections in a non-manifold grid

The defining feature of FoamGrid is its capability to handle network grids, i.e., grids with a non-manifold topology (Figures 1 and 2). In particular, more than two elements can meet in a common vertex in a one-dimensional grid, and more than two triangles can meet in a common edge in a two-dimensional grid.

Information about how a given element relates to its neighbors is essential for finite volume and DG methods, or more generally all methods involving element boundary fluxes. For this, the Dune grid interface provides the notion of *intersections*. In Dune terminology, an *intersection* is the set-theoretic intersection between the closures of two neighboring elements. Only intersections that have positive $(d-1)$–dimensional measure qualify as intersections in the Dune sense, and those are equipped with a coordinate system, i.e., a map from a $(d-1)$–dimensional reference element $T_{\text{ref,i}}$ to the intersection (Figure 3). Note that if the grid is conforming, then intersections will be geometrically the same as shared element faces. However, in general non-conforming grids, intersections are only subsets of element faces.

Intersections provide all information needed to compute element boundary fluxes. In particular, for each point on an intersection, you can get the vector $\mathbf{n}$ that is tangent to the element at that point, and normal to the element boundary. Also, you get the embeddings of the intersection into the two elements, in form of maps from the intersection reference element $T_{\text{ref,i}}$ to the element reference elements $T_{\text{ref,tr}}$ (Figure 3). Those maps pick up the same ideas used to model element geometries in Section 2.1. They are called *geometry-in-inside* and *geometry-in-outside*, respectively.

In C++ code, *intersections* are objects of type `Intersection`. For any given element (which is then called the *inside* element), all its intersections with neighboring elements can be accessed by traversing them with a dedicated iterator. Using range-based `for` syntax, a loop over all intersections of the element called `element` is written as

```
1 for (const auto& intersection : intersections(foamGridLeafView,element))
2 {
```

```
3     // do something with 'intersection'
4   }
```

See the `dune-grid` class documentation for the details on the user interface of the `Intersection` class.

Unfortunately, the Dune intersection mechanism, as flexible as it is, is not flexible enough for network grids. The original idea was that while elements may have more than one neighbor across a given facet, there is (even in non-conforming situations) at most one neighbor *at any given point* on that facet. This reflects the assumption that computational domains are expected to be topological manifolds. At the time of writing, this assumption is still reflected in the grid interface. In particular, the method

```
1 bool neighbor() const;
```

(a public member of the `Intersection` class), informs whether there is a neighbor across a given intersection. However, this information is insufficient in network grids, where even in a conforming grid there may be an arbitrary number of neighbors meeting at a common intersection. Finite volume and DG methods need access to this group of neighbors, to know how to distribute the flux across this intersection.

For this reason the intersection concept and programmer interface is being revisited, and is likely to undergo changes in the future to better support network grids. Changes to the Dune interface can only be made in a democratic process, and it is therefore unclear at what point in time such a change will happen.

Two different approaches have been proposed to the Dune grid development community. We describe them both briefly here. The full proposal text is available at `www.dune-project.org/modules/dune-foamgrid`.

The first approach tries to be as minimally invasive as possible. In particular, it retains the notion of an intersection as an object that relates *two* elements. The extension consists of two semantic rules, and a change to the `neighbor` method. The first of the semantic rules makes sure that the "number of neighbors" across a given intersection is well-defined. Remember that the *geometry-in-inside* is, roughly speaking, the intersection interpreted as a subset of the element T_1.

Rule 1 *For any two intersections of a given element T_1, the geometries-in-inside are either disjoint or identical.*

With the number of neighbors properly defined, the `neighbor` method is generalized to return this number instead of a yes/no answer:

```
1 std::size_t neighbor() const;
```

Note that this change is fully backward-compatible, as intersections in a non-network grid will return either `1` or `0` here, which casts to the values `true` or `false` as used previously.

The second semantic rule provides a way to find all intersections that share a common geometry-in-inside. No additional interface method is added for this. Rather, it is guaranteed that all such intersections appear consecutively when traversing the intersections with the intersection iterator.

Rule 2 *If more than one neighbor is reachable over a given geometry-in-inside, then all intersections for this geometry-in-inside shall be traversed consecutively by the intersection iterator.*

With this rule, groups of neighbors can be identified and used in flux computations.

The second approach is more radical, because it changes the idea of an intersection itself. Intersections cease to be objects that relate *pairs* of elements. Rather, they now become objects

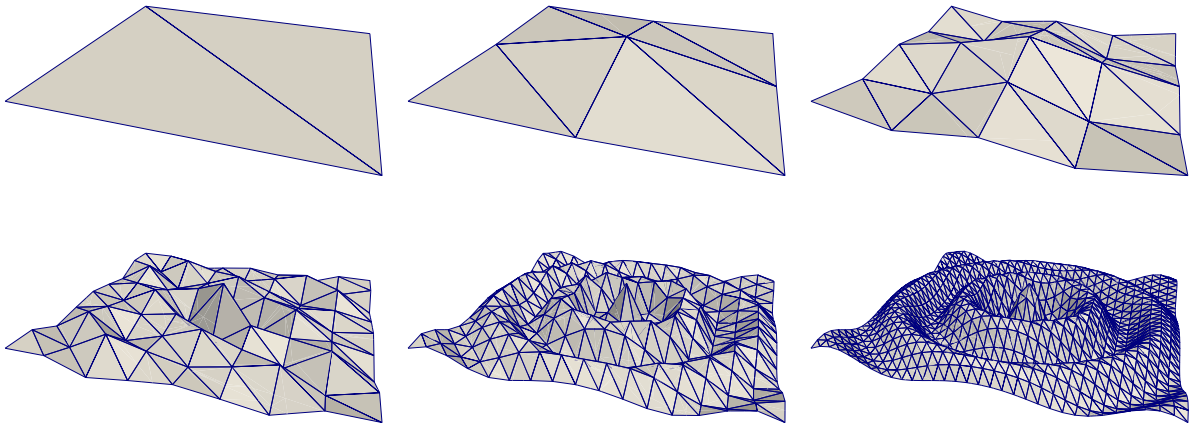

Figure 4: Grid refinement with element parametrizations

that relate *groups* of elements. As a consequence, each intersection still has only one geometry-in-inside. However, for each intersection there is now more than one outside element, each with corresponding geometry-in-outside and index-in-outside.

To access this information, more interface methods need to be changed. First of all, the `neighbor` method needs to be changed as proposed above. Secondly the methods

```
Entity outside () const;
LocalGeometry geometryInOutside () const;
int indexInOutside () const;
```

of the `Intersection` interface class need to be replaced by

```
Entity outside (std::size_t i=0) const;
LocalGeometry geometryInOutside (std::size_t i=0) const;
int indexInOutside (std::size_t i=0) const;
```

respectively. Rather than returning the unique outside element or its geometry or index, the methods must now return the corresponding quantity for the `i`-th outside element.

These changes are again fully backward-compatible. In grids without multiple intersections, at most the 0-th outside element will be available. The default parameter ensures that this intersection will be returned when the method is called without argument.

As an advantage, this proposal retains the rule that the geometries-in-inside must form a disjoint partition of the element boundary (modulo zero-sets). Also, it is easier to attach data to such intersections, which is a feature that has been requested various times in the past. On the downside, to iterate over all neighbors of an element, two nested loops are needed, instead of only one. This will make some code a bit longer, and more difficult to read.

Both proposals are currently under discussion. However, even with the current status quo, applications involving fluxes can be written for network domains. One-dimensional domains are straightforward as there the intersections are a fortiori conforming (see Sections 4.2 and 4.3). Two-dimensional networks need more trickery, but can also be made to work. Once either of the proposed interface extensions has been officially accepted, implementations of such methods will be much simpler.

3.2 Element parametrizations

In a standard non-conforming red refinement algorithm, new vertices are placed at the edge midpoints of refined elements. That way, while the grid gets finer and finer, the geometry of the

grid remains identical to the geometry of the coarsest grid. To also allow improvements to the geometry approximation, FoamGrid can use element parametrizations. Each coarse grid element T with reference element T_{ref} of a FoamGrid can be given a map

$$\varphi_T : T_{\text{ref}} \to \mathbb{R}^w,$$

which describes an embedding of T into physical space $\mathbb{R}^w$. This does not influence the grid itself — elements of a FoamGrid are always affine. However, when refining the grid, new elements are not inserted at edge midpoints. Rather, for each new vertex which would appear at an edge midpoint in standard refinement, the corresponding local position in its coarsest ancestor element T is determined using the refinement tree. This position is then used as an argument for φ_T, and the result is used as the position of the new vertex. That way, the grid approaches the shape described by the parametrization functions φ_T more and more as the grid gets refined (Figure 4).

In `dune-grid`, element parametrizations are implemented as small C++ objects. These must inherit from the abstract base class

```
VirtualFunction<FieldVector<double,dim>,
                FieldVector<double,dimworld> >;
```

declared in `dune/common/function.hh`.[4] One object of this type needs to be created for each element of the coarse grid. The base class has a single pure virtual method

```
virtual void evaluate(const FieldVector<double,dim>& x,
                      FieldVector<double,dimworld>& y) const = 0;
```

which implements the evaluation of φ_T. The first argument `x` is a position in local coordinates of the corresponding coarse grid triangle. The result $\varphi_T(x)$ is returned in the second argument `y`.

Element parametrizations are handed to the `GridFactory` during grid construction. Normally, grid elements are entered using the method

```
void insertElement(const GeometryType& type,
                   const std::vector<unsigned int>& vertices);
```

of the `GridFactory` class. For parametrized elements, there is the alternative method

```
void insertElement(const GeometryType& type,
                   const std::vector<unsigned int>& vertices,
                   const std::shared_ptr<VirtualFunction<
                                           FieldVector<ctype,dim>,
                                           FieldVector<ctype,dimworld>
                                         > >& elementParametrization);
```

This inserts the element with vertex numbers given in the array `vertices` and an element parametrization given by the object `elementParametrization` into the grid. The `GmshReader` uses this method for some of the higher-order elements that can appear in GMSH grid files.

To see the effect of element parametrizations, Figure 4 shows an example where the coarsest grid consists of only two triangles covering the domain $\Omega = [-1,1]^2 \times \{0\}$. As a parametrization we use the global function

$$\varphi : \mathbb{R}^2 \to \mathbb{R}^3, \qquad \varphi\begin{pmatrix} x_1 \\ x_2 \end{pmatrix} = \begin{pmatrix} x_1 \\ x_2 \\ 0.2 \cdot \exp(-|x|)\cos(4.5\pi|x|) \end{pmatrix}, \tag{1}$$

[4] For various reasons there is discontent with using the `VirtualFunction` class to implement element parametrizations. Therefore, it is expected to be replaced by something else eventually. However, at the time of writing, there is no specific proposal for such a replacement. In any case, while the details of the implementation are under discussion, the general idea of element parametrizations is not disputed, and can be expected to remain the way it is described here.

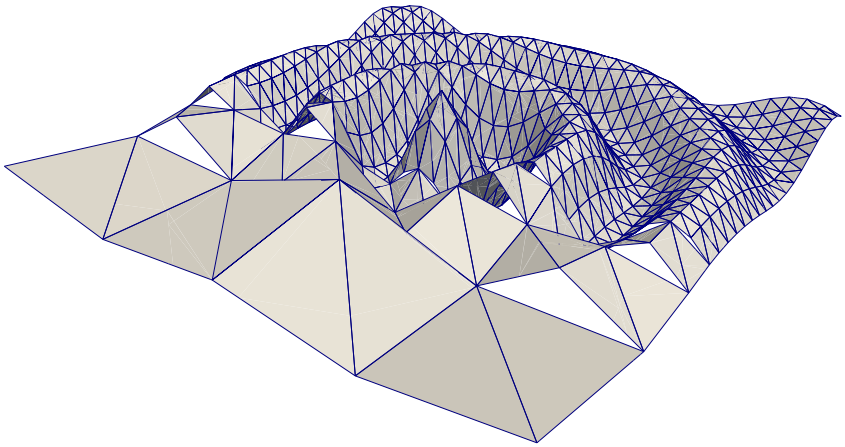

Figure 5: Adaptive refinement with element parameterizations leads to non-conforming geometries

and each element parametrization φ_T first maps its local coordinates to $\mathbb{R}^2$, and then applies the global function φ. Hence, upon refinement, the grid will approach the graph of the function (1). The code for this example is provided in the `dune-foamgrid` module itself, in the file `dune-foamgrid/examples/parametrized-refinement.cc`.

There is one pitfall when using element parametrizations together with non-conforming adaptive refinement for two-dimensional grids. If a hanging node appears in the course of refinement, the position of this node is determined by the parametrization functions of the corresponding element, which is typically *not* the midpoint of the adjacent longer edge. As a consequence, the grid will have holes wherever different refinement levels meet (Figure 5). While this may seem surprising at first sight, it is nevertheless a logical consequence of the hanging nodes refinement. The continuous domain is approximated by discontinuous grids, in direct correspondence to how discontinuous FE functions may be used to approximate continuous PDE solutions.

The non-conforming geometry approximation shows another shortcoming of the current Dune intersection concept. There, intersections are defined as set-theoretic intersections of (closures of) neighboring elements. However, in the geometrically non-conforming situation, neighboring elements do not actually intersect, and one can only speak of logical intersections. The practical consequence of this is that intersections do not have a single well-defined shape in $\mathbb{R}^w$ anymore. Rather, there are now *two* of them. This possibility to have two global element shapes is not currently reflected by the Dune grid interface. In the current implementation of FoamGrid, the method `Intersection::Geometry` will always return the global shape of the intersection as seen from the inside element.

3.3 Grid growth

A certain number of network problems is posed on domains that grow and/or shrink in the course of the simulation. Examples are fracture growth processes and simulations of bone trabeculae remodeling. To support such simulations, FoamGrid objects are allowed to grow and shrink, i.e., elements can be added and removed from the grid at runtime. Finite element and finite volume data is kept during such grid changes, using an approach much like the one used for grid adaptivity. Of all Dune grids, FoamGrid is currently the only one to support this feature. The programmer interface combines ideas from the `GridFactory` class to insert new vertices and elements, with the adaptivity interface to allow to keep data across steps of grid growth.

Growth and shrinkage of grids is a two-step process. First, new elements and vertices are

handed to the `FoamGrid` object. These are not inserted directly; rather, they are queued for eventual insertion. In addition, individual elements can be marked for removal. Once all desired elements and removal marks are known to the grid, the actual grid modification takes place in a second step.

Queuing elements for insertion and removal is controlled by three methods. The first,

```
std::size_t insertVertex(const FieldVector<double,dimworld>& x);
```

queues a new vertex with coordinates `x` for insertion. The return value of the method is an index that can be used to refer to this vertex when inserting new elements. The index remains fixed until all queued elements are actually inserted in the grid (by the `grow` method), but may change during the execution of that method. Inserting elements is done using

```
void insertElement(const GeometryType& type,
                   const std::vector<unsigned int>& vertices);
```

which mimics the corresponding method from the `GridFactory` class. The argument `type` has to be a simplex type, because (currently) FoamGrid supports only simplex elements. The array `vertices` must contain the indices of the vertices of the new element to be inserted. These can be either indices of existing vertices, or new indices obtained as the return values of the `insertVertex` method.

Analogously a new element with a parametrization can be inserted by calling

```
void insertElement(const GeometryType& type,
                   const std::vector<unsigned int>& vertices,
                   const std::shared_ptr<VirtualFunction<
                                            FieldVector<ctype,dim>,
                                            FieldVector<ctype,dimworld>
                                         > >& elementParametrization);
```

Finally, the method

```
void removeElement(const Codim<0>::Entity& element);
```

marks the given element for removal.

Once all desired elements are queued for insertion or removal, the actual grid modification takes place in a second step. The grid is modified using the method

```
bool elementsInserted = grid->grow(); // true if at least one element
    was inserted
```

While element removal is guaranteed, queuing elements does not assure that the element will be inserted. New elements are restricted by the fact that Dune grids are hierarchic objects. The vertices given by the user to form an element are always leaf vertices but may be contained in different hierarchic levels. However, elements can only be constituted by vertices of the same level. Therefore, new elements in FoamGrid are always inserted on the lowest possible level substituting the given vertex by its hierarchic descendants or ancestors. Note that it is not generally guaranteed that relatives of the given vertices on the same level can be found. In that case, the element will not be inserted. The method `grow` will return `true` if it was possible to insert at least one element.

After the call to `grow`, it is possible to check whether a given element has been created by the last call to the `grow` method:

```
bool isNew = element.isNew();     // true if element was created by last
    growth step
```

which is a method of the interface class `Entity<0>`, i.e. elements. Observe that this is the same method that returns whether an element has been created by grid refinement. Hence its semantics depends on whether it is queried after a call to `grow` or after a call to `adapt`. Using this method is helpful, e.g., when setting initial values and/or boundary conditions for newly created elements.

The growth is completed with the call

```
1 grid->postGrow();
```

which removes all `isNew` markers.

Summing up, growing or shrinking the grid while keeping grid data consists of the following steps. Note the relationship to grid adaptivity with data transfer.

1. Mark elements for removal; queue new vertices and elements for insertion.
2. Transfer all simulation data attached to the grid into an associative container indexed by the entity ids described in Section 2.2.
3. Call `grow()`.
4. Resize data array; copy data from the associative container into the array.
5. Set initial data at newly created elements and vertices and boundary conditions at newly formed boundaries. Note that element removal always creates new boundaries.
6. Finalize by calling `postGrow()`.

While grid growth itself is straightforward, it is difficult to use in combination with adaptive refinement. Grid refinement in Dune leads to hierarchical grids, which are forests of refinement trees. Not every element of such a forest can be added or removed without violating certain consistency conditions. In one-dimensional grids, there are relatively few problems, and grid growth and refinement can be used together to good effect. For two-dimensional grids we have tried to be as general as possible, but there are limits.

There are obstacles both to the removal of elements and to the insertion of new ones, if grid refinement is involved. First of all, only leaf grid elements can be removed. There is no conceptual problem with element removal if the element has no father. If it does have a father, however, removing only a subset of its 2^d sons is a violation of the Dune grid interface specification, which mandates that the sons of an element must (logically) cover their father [Bastian et al., 2008b, Def. 11.1]. Deliberately allowing this violation nevertheless appears to be the only viable solution, as all other possibilities amount to only allowing the removal of large groups of elements, which is rarely desired.[5]

Inserting new elements into a refinement forest of elements is even more problematic, because the new element needs to be assigned a level number in the hierarchy. Ideally, this level would always be zero, because if the element had a larger level number it would be expected to have a father. FoamGrid does violate this assumption if necessary, which does not lead to problems in practice, unless multigrid-type algorithms are used. On the other hand, the element level must be the same as the levels of all of its vertices. If the vertices are new, their levels can be freely chosen. However, grid growth almost always involves vertices that already exist in the grid. When trying to insert elements with vertices having different level numbers, FoamGrid currently tries to replace existing vertices by their father vertices, to obtain a set of $d + 1$ vertices on the same level (which then determines the element level).

4 Numerical examples

We close the article with three numerical examples. These show how seemingly challenging algorithms can be implemented with ease using the FoamGrid grid manager.

[5] Another Dune module which violates this assumption to good effect is `dune-subgrid`, see [Gräser and Sander, 2009].

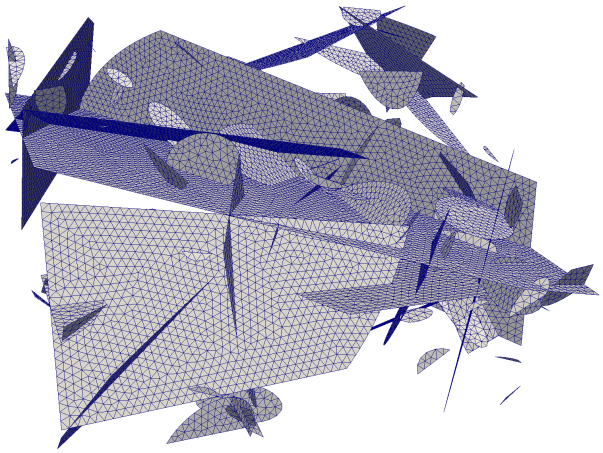

Figure 6: Grid for a discrete fracture network, courtesy of Patrick Laug. The network and grid were generated using the software described in [Borouchaki et al., 2000].

4.1 Unsaturated Darcy flow in a discrete fracture network

For our first example we simulate unsaturated Darcy flow through a network of two-dimensional fractures embedded into $\mathbb{R}^3$. We only consider the flow in the network itself, but FoamGrid can be easily coupled to higher-dimensional background grids using the `dune-grid-glue` module [Bastian et al., 2010, Engwer and Müthing, to appear].

Let the computational domain Ω be the union of a finite number of closed, bounded hypersurfaces in $\mathbb{R}^3$. We use the Richards equation to model the flow in Ω [Bear, 1988]. That is, we suppose that the state of the system in a time interval $[0, T]$ can be described by single scalar pressure field

$$p : \Omega \times [0, T] \to \mathbb{R}.$$

From this pressure field, the fluid saturation θ and permeability kr can be computed using the Brooks–Corey and Burdine parameter functions

$$\theta(p) = \begin{cases} \theta_m + (\theta_M - \theta_m) \left(\frac{p}{p_b}\right)^{-\lambda} & \text{for } p \le p_b \\ \theta_M & \text{for } p \ge p_b, \end{cases} \qquad \mathrm{kr}(\theta) = \left(\frac{\theta - \theta_m}{\theta_M - \theta_m}\right)^{3+\frac{2}{\lambda}},$$

where θ_m, θ_M, p_b, and λ are scalar parameters. The saturation θ satisfies the Richards equation

$$\frac{\partial}{\partial t}\theta(p) + \operatorname{div} \mathbf{v}(x, p) = 0, \qquad \mathbf{v}(x, p) = -K(x)\,\mathrm{kr}(\theta(p))\nabla(p - \varrho g z). \tag{2}$$

For simplicity we suppose that the flow is purely driven by the boundary conditions, and omit the gravity term $\varrho g z$.

Equation (2) is a quasilinear equation in the pressure p. In [Alt and Luckhaus, 1983] (see also [Berninger et al., 2011]) it was shown how the Kirchhoff transformation can be used to transform it to a semilinear equation for a generalized pressure

$$u : \Omega \times [0, T] \to \mathbb{R}, \qquad u(x, t) = u(p(x, t)) := \int_0^p \mathrm{kr}(\theta(q(x, t)))\, dq.$$

We discretize this equation in time using an implicit Euler method, and in space using first-order Lagrangian finite elements. The resulting weak discrete spatial problem can be written as a minimization problem for a strictly convex functional. At each time step, we determine the minimizer of this functional by a monotone multigrid method [Berninger et al., 2011].

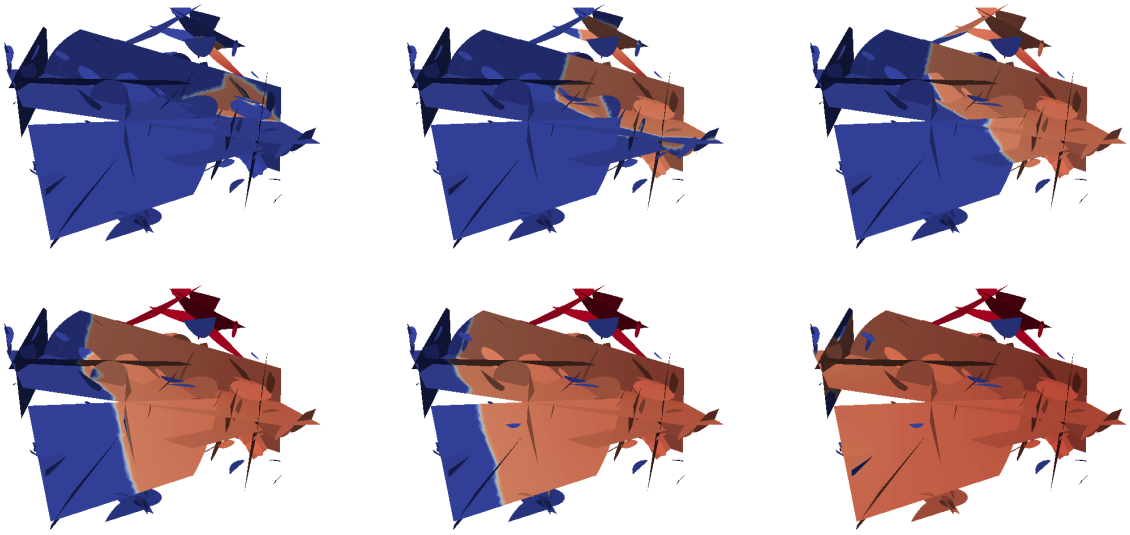

Figure 7: Unsaturated Darcy flow in a fracture network. The color visualizes the physical pressure p.

For the implementation we used the code used for the numerical examples in [Berninger et al., 2011]. Since vertex-based finite elements were used for the discretization, no changes to the numerical algorithm were needed to also apply it to network grids. Originally, the code used the Dune libraries and the UGGrid grid manager for unstructured grids. Even though the code for [Berninger et al., 2011] was not written with network flow problems in mind at all, it could nevertheless be reused as is, after only a handful of trivial bugfixes. The only changes necessary were replacing the line

```
typedef UGGrid<2> GridType;
```

by

```
typedef FoamGrid<2,3> GridType;
```

and adjusting the boundary data specification. This once more proves the point that using the Dune grid interface gives great flexibility, and allows code to do more things than planned by the original authors.

We present an example simulation using an artificially created network grid. The grid, which can be seen in Figure 6, was created by Patrick Laug using the fracture network grid generator described in [Borouchaki et al., 2000]. The grid spans the volume $[-6.5, 6.5]^2 \times [-2.165, 2.165]$ (length and pressure are given in meters). We assume the network to be filled with a sand-like material with parameters $\theta_m = 0.0458$, $\theta_M = 1$, $K = 6.54 \cdot 10^{-5}\,\mathrm{m/s}$, bubbling pressure $p_b = 0.0726\,\mathrm{m}$, and $\lambda = 0.694$.

We assume the network to be initially devoid of water, setting $p = -10\,\mathrm{m}$. Then, water is injected in a unit circle centered at the point $(0, 6.5, 0)$ on the boundary $\partial\Omega \cap \{x_1 = 6.5\}$ with a constant pressure of $p = 3\,\mathrm{m}$; no-flow boundary conditions are imposed at the remaining boundary. Figure 7 shows several steps in the evolution of the physical pressure p. As expected, one can see the fluid entering, and slowly filling almost the entire network. At the end, a steady state is reached, and the pressure is constant in each connected component of the domain.

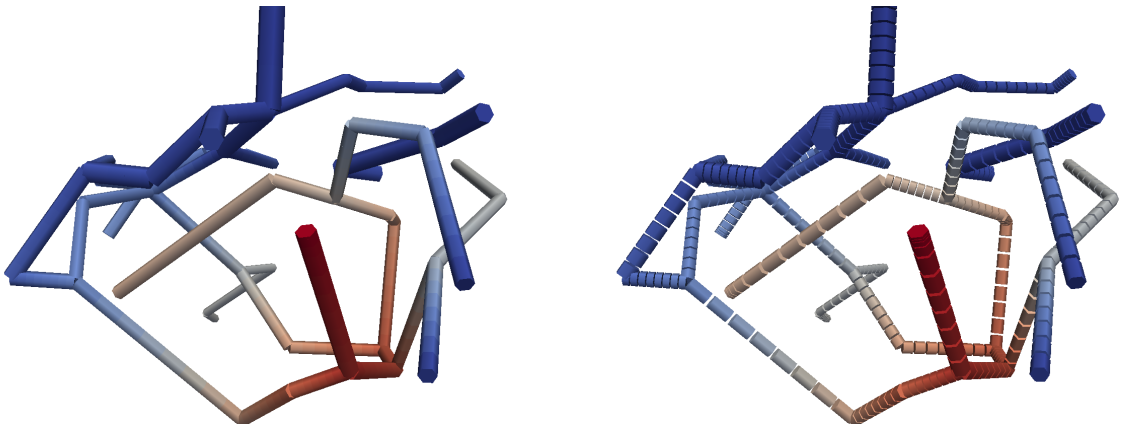

Figure 8: Single-phase two-component flow in a blood vessel network. One-dimensional elements are rendered as tubes. The color visualizes the fluid pressure. Left: stationary pressure p; right: initial element sizes.

4.2 Transport of a therapeutic agent in the microvasculature

The next example demonstrates local grid adaptivity, and the treatment of network bifurcations. We model the propagation of a therapeutic agent for cancer therapy in a network of small blood vessels. To reduce the computational effort, the network of three-dimensional vessels is reduced to a one-dimensional network Ω embedded in a three-dimensional tissue domain $\mathcal{T}$.

We first discuss the one-dimensional model for flow in a blood vessel segment, disregarding any bifurcations. Starting from a three-dimensional blood vessel domain Ξ, we describe blood in the microvasculature as an incompressible Newtonian fluid with viscosity μ and density ϱ governed by the Stokes equation. Further, we assume an axially symmetric blood vessel segment with constant radius R, cross-section area A, the parametrized tangent on the vessel centerline $\lambda(s)$, $\lambda : \mathbb{R} \to \mathbb{R}^3$, $s \mapsto \lambda(s)$, and a rigid vessel wall. Then, with the assumption of constant pressure p in a cross-section and negligible radial velocities v_r the Stokes equations can be integrated over a vessel segment yielding the one-dimensional equation

$$\left.\begin{aligned} \frac{A}{\varrho}\frac{\partial p}{\partial s}\lambda + 2\pi\frac{\mu}{\varrho}(2+\gamma)\bar{v}\lambda = 0 \\ \frac{\partial}{\partial s}\left(\frac{\pi R^4}{2\mu(2+\gamma)}\frac{\partial p}{\partial s}\right) = 0 \end{aligned}\right\} \quad \text{in } \Omega. \tag{3}$$

The parameter γ shapes a power-type axial velocity profile with mean velocity $\bar{v}$, yielding a quadratic velocity profile for $\gamma = 2$ and flatter profiles for $\gamma > 2$. For a detailed derivation in a more general setting, see the reduction of the Navier–Stokes equations to one-dimensional equations in [Quarteroni and Formaggia, 2003]. Equation (3) is a simplified stationary version of the one-dimensional blood flow equations described by Quarteroni and Formaggia [2003], neglecting vessel wall displacement and intertial forces.

Small blood vessels are exchanging mass with the embedding tissue through the vessel wall. The fluid exchange can be modeled by Starling's law, and results in an additional source term. With this modification, (3) becomes

$$\left.\begin{aligned} \frac{A}{\varrho}\frac{\partial p}{\partial s}\lambda + 2\pi\frac{\mu}{\varrho}(2+\gamma)\bar{v}\lambda = 0 \\ \frac{\partial}{\partial s}\left(\frac{\pi R^4}{2\mu(2+\gamma)}\frac{\partial p}{\partial s}\right) - 2\pi R L_p(p-\bar{p}_i) = 0 \end{aligned}\right\} \quad \text{in } \Omega, \tag{4}$$

where L_p is the empirical filtration coefficient dependent on, e.g., the intrinsic permeability and thickness of the membrane, and the viscosity of the interstitial fluid. The source term further depends on the pressure in the surrounding tissue $\bar{p}_i$. For the sake of simplicity, this tissue pressure is subsequently assumed constant. Equation (4) was also used by Cattaneo and Zunino [2014a] in a finite element setting to model coupled vessel-tissue flow processes in a tumor tissue.

The transport of a therapeutic agent is modeled by an advection–diffusion equation using the velocity field calculated by equation (4). Similar to the reduction of the Stokes equations we can reduce the three-dimensional advection–diffusion equation by integration over a vessel segment assuming a constant concentration $c = x\varrho$ on a given cross-section with area $A = \pi R^2$ [D'Angelo, 2007, Cattaneo and Zunino, 2014b]. The transport over the vessel wall can be described by the Kedem–Katchalsky equation [Kedem and Katchalsky, 1958], yielding

$$\frac{\partial(Ac)}{\partial t} + \bar{v}\frac{\partial(Ac)}{\partial s} - D_e\frac{\partial^2(Ac)}{\partial s^2} - 2\pi R\left[L_c(c-\bar{c}_i) + L_p(p-\bar{p}_i)(1-\sigma_c)c\right] = 0 \quad \text{in } \Omega. \tag{5}$$

The last term in (5), accounting for transport across the vessel wall, consists of an advective and a diffusive part where advection is reduced by the reflection coefficient $\sigma_c \in [0,1]$ for larger molecules. Again, we assume the mean tissue concentration $\bar{c}_i$ to be constant.

To model a network of such segments we split the blood vessel network Ω at junctions into pieces yielding a set of vessel segments Ω_i each governed by equations (4) and (5). At each junction we require continuity of pressure

$$p = p_1 = ... = p_j.$$

Together with these coupling conditions, and boundary conditions on $\partial\Omega$, Equation (4) has a unique solution, which is the stationary pressure field p.

For the transport at the junctions, we require continuity of concentration

$$c = c_1 = ... = c_j,$$

and, for Q_i, the volume flux leaving segment Ω_i at the junction, we require mass conservation

$$\sum_{i=1}^{j} Q_i = 0.$$

We discretize equations (4) and (5) in space with a standard finite volume scheme and piecewise linear one-dimensional grid elements. This demands balancing fluxes over the edges of the elements. Assume that we have calculated all element transmissibilities

$$t_i \coloneqq \frac{\pi R_i^4}{2\mu(2+\gamma)}. \tag{6}$$

Then, the volume flux Q_{ij} from element i to a neighboring element j can be calculated as

$$Q_{ij} = t_{ij}(p_i - p_j) = \frac{t_i t_j}{\sum_{k=0}^{N} t_k}(p_i - p_j). \tag{7}$$

One can see that the transmissibility at branching points is dependent on the transmissibility of all N neighboring elements. Assuming that all element transmissibilities t_i reside in an array `transmissibility`, the calculation of the two-point transmissibilities t_{ij} could look as follows using FoamGrid. Note that only a single loop over the grid entities is necessary to calculate the transmissibilities.

```
// for each element with index eIdx
std::vector<double> tSums(e.subEntities(/*codim=*/ 1), 0.0);
std::vector<double> tij;
std::vector<std::size_t> neighborFacetMap;

// loop over all intersections of this element
for(const auto& intersection : intersections(foamGridLeafView, element))
{
  if(intersection.neighbor())
  {
    std::size_t nIdx =
       foamGridLeafView.indexSet().index(intersection.outside());
    tij.push_back(transmissibility[eIdx]*transmissibility[nIdx]);
    neighborFacetMap.push_back(intersection.indexInInside());
    tSums[intersection.indexInInside()] += transmissibility[eIdx];
  }
  if(intersection.boundary())
    // boundary treatment ...
}

// compute the two-point transmissibilities
for (std::size_t i = 0; i < tij.size(); i++)
  transmissibilitiesIJ[i] /= tSums[neighborFacetMap[i]];
```

This code works well using the grid interface of `dune-grid-2.4`, even though that interface does not have any provisions for network grids at all (Section 3.1). This is because the grid is one-dimensional. In this case, each intersection always covers entire element facets (i.e., vertices). Therefore, the `indexInInside` method can be used to group intersections.

To reduce numerical diffusion induced by the implicit Euler scheme, the grid can be refined adaptively around the concentration front. Initially, the grid is refined around inflow boundaries, and during the simulation the grid is adapted by a local gradient-based concentration indicator as described in [Wolff, 2013]. This indicator marks an element i for refinement if

$$\frac{\max_i(\Delta c_{ij}) - \Delta c_{\min}}{\Delta c_{\max} - \Delta c_{\min}} \geq \epsilon_r,$$

and for coarsening if

$$\frac{\max_i(\Delta c_{ij}) - \Delta c_{\min}}{\Delta c_{\max} - \Delta c_{\min}} < \epsilon_c,$$

where Δc_{ij} denotes the concentration difference between element i and a neighboring element j. The parameters ϵ_r, $\epsilon_c \in [0, 1]$ ($\epsilon_r > \epsilon_c$) are problem dependent. We achieved robust adaptation behavior in our example for $\epsilon_r = 0.3$ and $\epsilon_c = 0.05$. The indicator is evaluated before every time step and marks elements for refinement or coarsening using the `mark` method, see Section 2.3. The adaption of the grid is then handled by FoamGrid. Between the `preAdapt` and `postAdapt` steps, we have to transfer data, i.e., primary variables and spatial parameters from the old grid to the new adapted grid. As can be seen in Figure 9, the refinement scheme works well even around network bifurcations.

We simulate a network of capillaries in rat brain tissue scanned by Motti et al. [1986] and reconstructed as segment network with three-dimensional geometrical information by Secomb et al. [2000]. The network data comprises three-dimensional location of vessel segments, inflow and outflow boundary markers, vessel segment radii, and velocity estimates for each vessel segment. The domain has a bounding box of $150\,\mu\text{m} \times 160\,\mu\text{m} \times 140\,\mu\text{m}$. The given vessel radii vary between $2\,\mu$m and $4.5\,\mu$m. The estimated velocities are in a range of $0.5\,\frac{\text{mm}}{\text{s}}$ to $7.5\frac{\text{mm}}{\text{s}}$. We choose the blood viscosity $\mu = 3.0Pa \cdot s$, the filtration coefficient $L_p = 3.33 \cdot 10^{-12}\frac{m}{Pa \cdot s}$, the effective diffusion

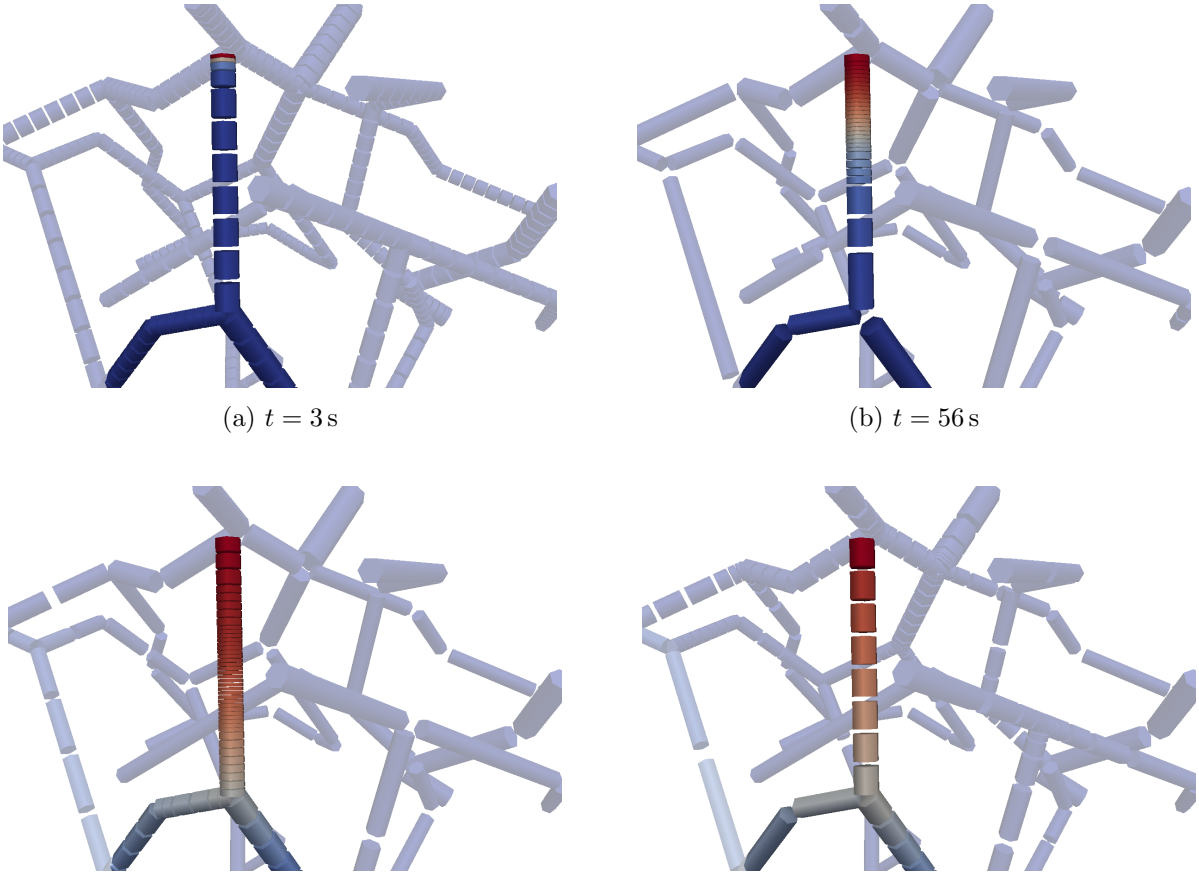

(a) $t = 3\,\mathrm{s}$ (b) $t = 56\,\mathrm{s}$

(c) $t = 201\,\mathrm{s}$ (d) $t = 3\,000\,\mathrm{s}$

Figure 9: Single-phase two-component flow in a blood vessel network. One-dimensional elements are rendered as tubes. The gaps merely exist to better visualize the adaptive grid. The images show the mole fraction x at different simulation times. Note how a concentration wave enters the network from the top, and is tracked by a zone of high grid resolution, even as it goes through a bifurcation point.

coefficient over the vessel wall $L_c = 10^{-5}\frac{1}{sm}$, and the diffusion coefficient of the transported agent trail $D_e = 2.93 \cdot 10^{-14}\,\frac{1}{\mathrm{s}}$, calculated with the Stokes–Einstein radius. We assign the following boundary conditions for the flow problem (4)

$$\begin{aligned} p &= p_D && \text{on } \partial\Omega_{\text{outflow}}, \\ \left[\frac{R^2}{\mu(2+\gamma)}\frac{\partial p}{\partial z}\right] \cdot n = \bar{v} \cdot n &= v_N && \text{on } \partial\Omega_{\text{inflow}}. \end{aligned}$$

The velocities for the inflow segments are set to the estimates provided by Secomb et al. [2000] along with the grid geometry. We simulate the arrival of a therapeutic agent by a Dirichlet boundary condition for equation (5) on a subset $\partial\Omega_{\mathrm{c}}$ of $\partial\Omega_{\text{inflow}}$, namely,

$$\begin{aligned} c &= c_D && \text{on } \partial\Omega_{\mathrm{c}}, \\ c &= 0 && \text{on } \partial\Omega_{\text{inflow}} \setminus \partial\Omega_{\mathrm{c}}. \end{aligned}$$

Specifically, we enforce a mole fraction of $x_D = 10^{-8}$ at one of the inflow vessel segments with a Dirichlet boundary condition. At outflow boundaries, we neglect diffusive fluxes. Note that a full upwind scheme is employed for the concentration, so no further boundary condition for the advective fluxes is necessary at outflow boundaries.

Figure 8 shows the resulting stationary pressure field and the initial element size. In Figure 9, one can see the resulting mole fraction at various times. Note how the local grid refinement follows the steepest gradients, i.e., the transport front. The resulting mole fractions vary from segment to segment due to different radii. This also automatically ensures a finer grid around vessel bifurcations. The one-dimensional elements are depicted as three-dimensional tubes scaled with their respective radius.

4.3 Root water uptake and root growth at plant-scale

In environmental and agricultural research fields, models describing root architectures are used to investigate water uptake and root growth behavior of plants [Dunbabin et al., 2013]. In our final example, a one-dimensional network embedded into $\mathbb{R}^3$ is used to describe such a plant root architecture. We simulate water flow through the root network and root growth.

Plant roots can be described by a tree-like network of pipes which consist of xylem tubes [Tyree and Zimmermann, 2002]. We follow the cohesion–tension theory [Tyree, 1997], where the water flow through the root system is governed by the pressure gradient caused by the transpiration rate of the plant above the soil. We assume only vertical flow and no gravity. This leads to a Darcy's law analogy [Doussan et al., 1998]

$$q_x = -K_x \frac{dp_x}{dz},$$

where q_x is the water flux in the xylem tubes, p_x is the xylem water pressure, and K_x is the axial conductance of one root segment.

Water can enter the roots at any point on the xylem tube surfaces, which leads to a volume source term for the one-dimensional network model. For simplicity, we model a single membrane only for the entire pathway of water from soil into the roots. With this assumption, and neglecting osmotic processes, radial water flow q_r into one root segment is defined as

$$q_r = K_r A_r (p_S - p_x),$$

where K_r is the radial conductivity, i.e., the conductivity of series of tissues from root surface to the xylem. The number $A_r := 2\pi r l$ is the soil–root interface area. The water pressure p_S at the soil–root interface must be provided. One option is to couple the root system to a Richards

equation based soil water flow simulation [Javaux et al., 2008], but for simplicity we simply take p_S as a known value.

The continuity equation leads to

$$-\operatorname{div}(q_x(p_x)) = S(p_x) \tag{8}$$

with a solution-dependent source

$$S(p_x) = K_r A_r (p_S - p_x),$$

where q_x is the only unknown variable. This modeling approach neglects the influence of solutes on water flow, as well as the capacitive effect of the roots, because the amount of water stored in roots is generally small compared to transpiration requirements.

Equation (8) is discretized in space with standard cell-centered finite volumes, piecewise linear one-dimensional grid elements and implemented using the external Dune discretization module DuMu^x [Flemisch et al., 2011], and the FoamGrid grid manager. Flux calculations over edges and branching points of the root network are implemented just as in our previous example.

So far, we have assumed a root network that does not alter its geometry over time. Several algorithms were developed to describe root growth (e.g., [Pagès et al., 2004, Somma et al., 1998, Leitner et al., 2010]). These models define root growth either by a fractal description, or more generically depending on plant specific parameters (root elongation, growth direction, branching density) and surrounding soil properties (soil moisture, soil strength, temperature, nutrients). For simplicity, we model root growth here as an (almost) completely random process. New root branches occur at random time steps and with a small gravity effect only, which makes the roots tend to point downwards. Existing branches grow at the branch tip at random times and without changing directions.

Our example simulation starts with a simple root grid which consist of one vertical root branch discretized with eight elements (root segments) and no lateral branches (Figure 10). We choose radial and axial conductivity values from [Doussan et al., 1998]. The surrounding relative soil pressure is set to $p_S = -2.9429 \cdot 10^{-2}$ Pa, and the Dirichlet boundary value at the root collar is set to $-1.2 \cdot 10^6$ Pa. Parameters and boundary conditions do not change with time.

In every time step, new root elements are created and either added to an existing lateral branch or to the main branch. An element-based indicator based on simulation time and random factors decides whether a new branch is added at one of the element's vertices. The `Indicator` class also computes the coordinates of the newly inserted point that is the second vertex of the new element. The coordinates depend again on simulation time, random factors, and the branch orientation in $\mathbb{R}^3$. Our growth step calculation in DuMu^x using FoamGrid looks as follows.

```
1  template<class Indicator>
2  void growGrid(Indicator& indicator, Variables& vars)
3  {
4    // (1) calculate indicator for each element
5    indicator.calculateIndicator();
6
7    // (2) insert elements according to the computed indicator
8    insertElements(indicator);
9
10   // (3) Put variables in a persistent map
11   storeVariables(vars);
12
13   // (4) Grow grid
14   grid->grow();
15
```

```
16    // (5) Resize and (re-)construction of variables
17    vars.resize(foamGridLeafView.size(0));
18    reconstructsVariables(vars);
19
20    // (6) delete isNew markers in grid
21    grid->postGrow();
22  }
```

The `vars` container contains all primary variables and spatial parameters defined on the root segments. The `Indicator` class is a template parameter that can be easily exchanged, to allow different root growth algorithms. In Step (2), the new elements are inserted. New root segments must be connected to the old grid. The implementation of the `insertElements` method could looks as follows:

```
1  template<class Indicator>
2  void insertElements(const Indicator& indicator)
3  {
4    // iterate over all grid elements (root segments)
5    for (const auto& element : elements(foamGridLeafView))
6    {
7      // find elements that will get a new neighbor
8      if (indicator.willGrow(element))
9      {
10       // get the new elements vertices from the indicator
11       std::size_t vIdx0 =
             grid->insertVertex(indicator.getNewVertexCoordinates(element));
12       // get index of the existing element vertex the new element will
             be connected to
13       std::size_t vIdx1 = indicator.getConnectedVertex(element);
14       // insert new element with the two vertex indices
15       grid->insertElement(element.type(), {vIdx0, vIdx1});
16     }
17   }
18 }
```

The primary variables and spatial parameters of our physical problem are stored in Step (3), before the actual growth step (4). We update the sizes of the variables and parameter vectors since the total number of degrees of freedom has changed due to the growth step (5). In addition, values for the new elements have to be computed. In our case, new root segments inherit the primary variables and the spatial parameters from its preceding neighbor element. Root tips, in particular new root tips, are always assigned Neumann no-flow boundary conditions. At the end, the `postGrow` method is called, which deletes the `isNew` markers (Step (6)).

Figure 10 shows the root network and the pressure distribution inside the roots for several time steps. The total root water uptake (transpiration demand) of the plant, defined by the Dirichlet boundary condition, does not change during the growing period. Thus, the pressure inside the plant changes since the water uptake of the plant is distributed to more and more roots segments.

References

H. Alt and S. Luckhaus. Quasilinear elliptic–parabolic differential equations. *Math. Z.*, 183: 311–341, 1983.

P. Bastian, M. Blatt, A. Dedner, C. Engwer, R. Klöfkorn, R. Kornhuber, M. Ohlberger, and O. Sander. A Generic Grid Interface for Parallel and Adaptive Scientific Computing. Part II: Implementation and Tests in DUNE. *Computing*, 82(2–3):121–138, 2008a.

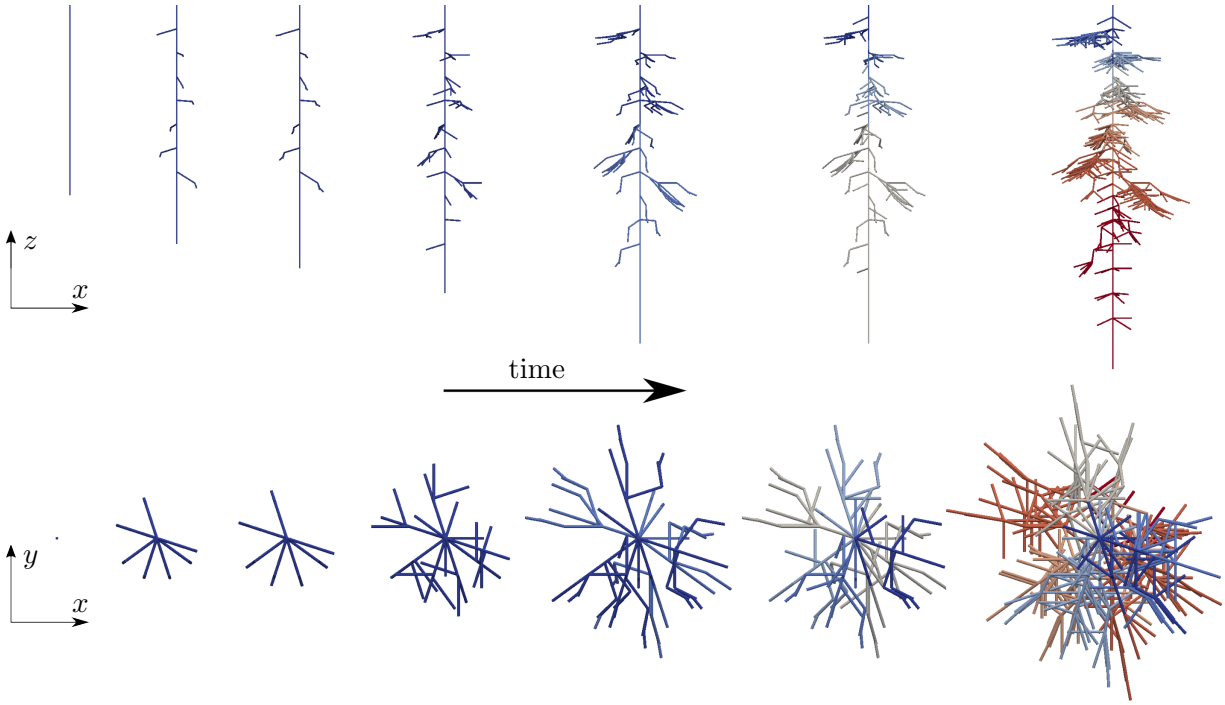

Figure 10: Growth of a root system, shown in a lateral (top) and an axial (bottom) view. The color represents the pressure inside the roots.

P. Bastian, M. Blatt, A. Dedner, C. Engwer, R. Klöfkorn, M. Ohlberger, and O. Sander. A Generic Grid Interface for Parallel and Adaptive Scientific Computing. Part I: Abstract Framework. *Computing*, 82(2–3):103–119, 2008b.

P. Bastian, G. Buse, and O. Sander. Infrastructure for the coupling of Dune grids. In *Proc. of ENUMATH 2009*, pages 107–114. Springer, 2010.

J. Bear. *Dynamics of Fluids in Porous Media*. Dover Publications, 1988.

H. Berninger, R. Kornhuber, and O. Sander. Fast and robust numerical solution of the Richards equation in homogeneous soil. *SINUM*, 49(6):2576–2597, 2011.

H. Borouchaki, P. Laug, and P. George. Parametric surface meshing using a combined advancing-front – generalized-Delaunay approach. *Int. J. Numer. Meth. Eng.*, 49(1–2):233–259, 2000.

A. Bressan, S. Čanić, M. Garavello, M. Herty, and B. Piccoli. Flows on networks: recent results and perspectives. *EMS Surv. Math. Sci.*, 1(1):47–111, 2014.

L. Cattaneo and P. Zunino. Computational models for fluid exchange between microcirculation and tissue interstitium. *Networks and Heterogeneous Media*, 9(1):135–159, 2014a.

L. Cattaneo and P. Zunino. A computational model of drug delivery through microcirculation to compare different tumor treatments. *International Journal for Numerical Methods in Biomedical Engineering*, 30(11):1347–1371, 2014b.

C. D'Apice, S. Göttlich, M. Herty, and B. Piccoli. *Modeling, simulation, and optimization of supply chains*. SIAM, 2010.

C. Doussan, L. Pages, and G. Vercambre. Modelling of the Hydraulic Architecture of Root Systems: An Integrated Approach to Water Absorption—Model Description. *Annals of Botany*, 81(2):213–223, 1998. doi: 10.1006/anbo.1997.0540.

V. M. Dunbabin, J. A. Postma, A. Schnepf, L. Pagès, M. Javaux, L. Wu, D. Leitner, Y. L. Chen, Z. Rengel, and A. J. Diggle. Modelling root–soil interactions using three-dimensional models of root growth, architecture and function. *Plant and Soil*, 372:93–124, 2013. doi: 10.1007/s11104-013-1769-y.

G. Dziuk and C. M. Elliott. Finite elements on evolving surfaces. *IMA J. Numer. Anal.*, 27: 262–292, 2007a.

G. Dziuk and C. M. Elliott. Surface finite elements for parabolic equations. *J. Comput. Math.*, 25(4):385–407, 2007b.

G. Dziuk and C. M. Elliott. Finite element methods for surface PDEs. *Acta Numerica*, 22: 289–396, 2013.

C. D'Angelo. Multiscale modelling of metabolism and transport phenomena in living tissues. *Bibliotheque de l'EPFL, Lausanne*, 2007.

C. Engwer and S. Müthing. Concepts for flexible parallel multi-domain simulations. In *Domain Decomposition Methods in Science and Engineering XXII*. Springer, to appear.

B. Flemisch, M. Darcis, K. Erbertseder, B. Faigle, A. Lauser, K. Mosthaf, P. Nuske, A. Tatomir, M. Wolff, and R. Helmig. DuMuX : DUNE for Multi-{ Phase, Component, Scale, Physics, ... } Flow and Transport in Porous Media. *Advances in Water Resources*, 34:1102–1112, 2011.

M. Garavello and B. Piccoli. *Traffic flow on networks*. American Institute of Mathematical Sciences (AIMS), 2006.

C. Geuzaine and F. Remacle. *Gmsh Reference Manual*, 2015. www.geuz.org/gmsh/doc/texinfo/gmsh.pdf.

C. Gräser and O. Sander. The dune-subgrid module and some applications. *Computing*, 8(4): 269–290, 2009.

S. Gross and A. Reusken. *Numerical Methods for Two-phase Incompressible Flows*. Springer, 2011.

M. Javaux, T. Schröder, J. Vanderborght, and H. Vereecken. Use of a Three-Dimensional Detailed Modeling Approach for Predicting Root Water Uptake. *Vadose Zone Journal*, 7(3):1079, 2008. doi: 10.2136/vzj2007.0115.

O. Kedem and A. Katchalsky. Thermodynamic analysis of the permeability of biological membranes to non-electrolytes. *Biochimica et biophysica Acta*, 27:229–246, 1958.

S. Lang, V. J. Dercksen, B. Sakmann, and M. Oberlaender. Simulation of signal flow in 3d reconstructions of an anatomically realistic neural network in rat vibrissal cortex. *Neural Networks*, 24(9):998–1011, 2011.

D. Leitner, S. Klepsch, G. Bodner, and A. Schnepf. A dynamic root system growth model based on L-Systems. *Plant and Soil*, 332:177–192, Jan. 2010. ISSN 0032-079X. doi: 10.1007/s11104-010-0284-7.

M. McClure, M. Babazadeh, S. Shiozawa, and J. Huang. Fully coupled hydromechanical simulation of hydraulic fracturing in three-dimensional discrete fracture networks. In *SPE Hydraulic Fracturing Technology Conference, 3–5 February, The Woodlands, Texas, USA*, 2015. doi: http://dx.doi.org/10.2118/173354-MS.

M. W. McClure and R. N. Horne. *Discrete Fracture Network Modeling of Hydraulic Stimulation – Coupling Flow and Geomechanics*. Springerbriefs in Earth Sciences. Springer, 2013.

E. D. Motti, H.-G. Imhof, and M. G. Yasargil. The terminal vascular bed in the superficial cortex of the rat: an SEM study of corrosion casts. *Journal of Neurosurgery*, 65(6):834–846, 1986.

S. Nammi, P. Myler, and G. Edwards. Finite element analysis of closed-cell aluminium foam under quasi-static loading. *Materials and Design*, 31:712–722, 2010.

I. Nitschke, A. Voigt, and J. Wensch. A finite element approach to incompressible two-phase flow on manifolds. *J. Fluid Mech.*, 708:418–438, 2012.

M. A. Olshanskii, A. Reusken, and J. Grande. A finite element method for elliptic equations on surfaces. *SIAM J. Numer. Anal.*, 47(5):3339–3358, 2009.

L. Pagès, G. Vercambre, and J. Drouet. Root Typ: a generic model to depict and analyse the root system architecture. *Plant and Soil*, 258:103–119, 2004.

A. Quarteroni and L. Formaggia. Mathematical Modelling and Numerical Simulation of the Cardiovascular System. In *Modelling of Living Systems*, Handbook of Numerical Analysis Series. EPFL, 2003.

S. Reuther and A. Voigt. The interplay of curvature and vortices in flow on curved surfaces. *SIAM Multiscale Model. Simul.*, 13(2):632–643, 2015.

M. E. Rognes, D. A. Ham, C. J. Cotter, and A. T. T. McRae. Automating the solution of PDEs on the sphere and other manifolds in FEniCS 1.2. *Geosci. Model Dev.*, 6:2099–2119, 2013.

A. Schmidt and K. G. Siebert. *Design of Adaptive Finite Element Software. The Finite Element Toolbox ALBERTA*, volume 42 of *LNCSE*. Springer, 2005. alberta-fem.de.

T. Secomb, R. Hsu, N. Beamer, and B. Coull. Theoretical simulation of oxygen transport to brain by networks of microvessels: effects of oxygen supply and demand on tissue hypoxia. *Microcirculation*, 7(4):237–247, 2000.

F. Somma, J. W. Hopmans, and V. Clausnitzer. Transient three-dimensional modeling of soil water and solute transport with simultaneous root growth, root water and nutrient uptake. *Plant and Soil*, 202:281–293, 1998.

A. V. T. Witkowski, R. Backofen. The influence of membrane bound proteins on phase separation and coarsening in cell membranes. *Phys. Chem. Chem. Phys.*, 14:14509–14515, 2012.

M. T. Tyree. The Cohesion-Tension theory of sap ascent : current controversies. *Journal of Experimental Botany*, 48(315):1753–1765, 1997.

M. T. Tyree and M. H. Zimmermann. *Xylem Structure and the Ascent of Sap*, volume 12. Springer, 2002. ISBN 9783540433545. doi: 10.1016/0378-1127(85)90081-7.

S. Vey and A. Voigt. AMDiS: adaptive multidimensional simulations. *Comput. Visual. Sci.*, 10: 57–67, 2007.

M. Wolff. *Multi-scale modeling of two-phase flow in porous media including capillary pressure effects*. PhD thesis, Universität Stuttgart, 2013.